Collaborative sparse regression using spatially correlated supports – Application to hyperspectral unmixing

Yoann Altmann, Marcelo Pereyra and Jose Bioucas Dias

arXiv:1409.8129v2 [stat.CO] 3 Oct 2015

Abstract

This paper presents a new Bayesian collaborative sparse regression method for linear unmixing of hyperspectral images. Our contribution is twofold; first, we propose a new Bayesian model for structured sparse regression in which the supports of the sparse abundance vectors are *a priori* spatially correlated across pixels (*i.e.*, materials are spatially organised rather than randomly distributed at a pixel level). This prior information is encoded in the model through a truncated multivariate Ising Markov random field, which also takes into consideration the facts that pixels cannot be empty (i.e, there is at least one material present in each pixel), and that different materials may exhibit different degrees of spatial regularity. Secondly, we propose an advanced Markov chain Monte Carlo algorithm to estimate the posterior probabilities that materials are present or absent in each pixel, and, conditionally to the maximum marginal *a posteriori* configuration of the support, compute the MMSE estimates of the abundance vectors. A remarkable property of this algorithm is that it self-adjusts the values of the parameters of the Markov random field, thus relieving practitioners from setting regularisation parameters by cross-validation. The performance of the proposed methodology is finally demonstrated through a series of experiments with synthetic and real data and comparisons with other algorithms from the literature.

Supported in part by the Direction Générale de l'armement, French Ministry of Defence, the EPSRC via grant EP/J015180/1, the SuSTaIN program - EPSRC grant EP/D063485/1 - at the Department of Mathematics, University of Bristol, by the Portuguese Science and Technology Foundation, Projects UID/EEA/50008/2013 and PTDC/EEI-PRO/1470/2012. M. Pereyra holds a Marie Curie Intra-European Fellowship for Career Development.

Y. Altmann is with the School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh U.K. (email: Y.Altmann@hw.ac.uk).

M. Pereyra is with the School of Mathematics of the University of Bristol, Bristol U.K. (email: marcelo.pereyra@bristol.ac.uk).

J.M. Bioucas-Dias is with the Instituto de Telecomunicações and Instituto Superior Técnico, Universidade de Lisboa, Portugal (email: bioucas@lx.it.pt).

Index Terms

Collaborative sparse regression, Spectral unmixing, Bayesian estimation, Markov random fields, Markov chain Monte Carlo methods.

I. INTRODUCTION

Spectral unmixing (SU) of hyperspectral images is a challenging problem that has received a lot of attention over the last few years [1]–[3]. It consists in identifying the materials (endmembers) present in an image and simultaneously quantifying their fractions or proportions within each pixel (abundances). This source separation problem has been widely studied for applications where pixel reflectances are linear combinations of pure component spectra [4]–[8]. The typical SU processing pipeline is then decomposed into three main estimation steps: the estimation of the number of different materials present the image, the estimation or extraction of their spectral signatures, and finally the quantification of their abundances.

Abundance estimation is known to be a challenging problem, in particular in scenarios involving materials with similar spectral signatures. In these cases, exploiting prior knowledge about the problem can improve estimation performance dramatically. A particularly important form of prior knowledge is that the number of materials within each pixel is typically much smaller than the total number of materials present in the scene (this property is accentuated in modern images that are acquired with high spatial resolution sensors). In other words, the abundance vectors are generally sparse. Sparsity also arises naturally when SU is performed with a dictionary or library containing the spectral signatures of a large number materials, some of which are possibly present in the scene [9].

Once sparsity is taken into consideration, SU can be conveniently formulated as a sparse regression problem whose objective is to jointly identify the materials within each pixel and to quantify their abundance. This regression problem is often solved by penalised maximum likelihood estimation, which can be efficiently computed with state-of-the-art optimisation algorithms (typically an ℓ_1 penalty is used to promote sparse solutions) [10]. Recently, Iordache *et al.* [11] proposed a collaborative sparse regression technique (CLSunSAL) based on an $\ell_{2,1}$ penalty function that enforces group-sparsity for the abundances of each material. This method was further improved in [12] by introducing a pre-processing step that identifies the elements from the spectral library that are more likely present in the image. The resulting MUSIC-CSR algorithm solves a sparse regression problem that is collaborative in the sense that all the image pixels are used to identify the active endmembers. Sparse regression for SU

can also be successfully performed within the Bayesian framework. For example, Dobigeon *et al.* [13] and Eches *et al.* [14] propose Bayesian models and Monte Carlo algorithms to identify the active endmembers in an HSI from a spectral library while ensuring that the abundances of absent endmembers are zero. Note that library-based methods are not the only strategy to address the absence of pure pixels (see [4], [15]–[17] for more details).

It is widely acknowledged that collaborative sparse regression methods can produce very accurate SU results. Collaboration is key because it improves the estimation of the support of the sparse abundance vectors, thus reducing significantly the number of unknowns. However, most existing collaborative techniques only exploit global information and therefore can only seek to determine if an endmember is present/absent in the entire image (that is, can only estimate the union of the supports of all the abundance vectors in the image). As a result, global collaborative techniques may overestimate significantly the support of the actual abundance vector of each pixel.

This paper presents a new collaborative sparse regression technique that exploits the local spatial correlations in the image to accurately detect the endmembers that are active/inactive in each pixel. Precisely, we present a Bayesian model that simultaneously promotes sparsity on the abundance vectors, and spatial correlation between the abundance vectors' supports (*i.e.*, non-zero elements), modelling the spatial presence and absence patterns of materials in the scene. This approach differs significantly from the strategies adopted in the previous works [18]–[22], where spatial correlations are introduced by regularising the abundance values or the nonlinear effects occurring in the image. The latter strategies perform well in hyperspectral images with low spatial resolution, and composed mainly of homogenous regions, but are inadequate for high resolution images and for images involving complex scenes, small targets, and textures or fluctuations in the abundances. By operating directly at the level of the abundance vectors' support, our model is able to capture spatial correlations in a more subtle manner and produce accurate estimation results in challenging scenes involving, for example, crops (textured scenes) and isolated trees.

The remainder of the paper is organized as follows. Section II recalls the classical linear mixing model for SU and presents the proposed Bayesian model for sparse regression. In Section IV we propose an original Monte Carlo method to perform Bayesian inference in this model and perform SU. The proposed methodology is demonstrated on synthetic and real HSI in in Sections V and VI. Conclusions are finally reported in Section VII.

II. Problem statement

Consider a hyperspectral image $\mathbf{Y} \in \boldsymbol{R}^{L\times N}$, where L is the number of spectral bands considered and $N = N_{\text{row}} \times N_{\text{col}}$ corresponds to the total number of pixels. Under the linear mixing assumption, each image pixel $\mathbf{y}_n = [\mathbf{Y}_{1,n}, \ldots, \mathbf{Y}_{L,n}]^T \in \boldsymbol{R}^L$ can be expressed as a linear combination of R known spectral signatures $\mathbf{m}_1, \ldots, \mathbf{m}_R$ corrupted by zero-mean Gaussian noise with diagonal covariance $\boldsymbol{\Sigma}_0$, that is,

$$\mathbf{y}_n = \sum_{r=1}^{R} a_{r,n}\mathbf{m}_r + \mathbf{e}_n, \quad \mathbf{e}_n \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_0) \tag{1}$$

where $a_{r,n}$ is the mixing coefficient associated with the rth endmember in the nth pixel and $\boldsymbol{\Sigma}_0 = \text{diag}(\boldsymbol{\sigma}^2)$ with $\boldsymbol{\sigma}^2 = [\sigma_1^2, \ldots, \sigma_L^2]^T$. By setting $\mathbf{M} = [\mathbf{m}_1, \ldots, \mathbf{m}_R]$ and $\boldsymbol{a}_n = [a_{1,n}, \ldots, a_{R,n}]^T$, Eq. (1) can be conveniently expressed in matrix notation as $\mathbf{y}_n = \mathbf{M}\boldsymbol{a}_n + \mathbf{e}_n$.

This paper considers the supervised spectral unmixing problem of the hyperspectral image $\mathbf{Y}$, *i.e.*, the estimation of the $R \times N$ abundance matrix $\mathbf{A} = [\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N]$. More precisely, we consider Bayesian methods for estimating $\mathbf{A}$ given $\mathbf{Y}$ (and the endmember matrix $\mathbf{M}$) subject to the following two sets of physical constraints: first, the abundances are non-negative quantities, *i.e.*, $a_{r,n} \geq 0 \quad \forall r, n$; second, there is at least one material present in each pixel and therefore $\|\boldsymbol{a}_n\|_0 > 0$, where $\|\cdot\|_0$ denotes the ℓ_0 vector pseudo-norm. We also assume that the values of the noise variances $\boldsymbol{\sigma}^2$ are unknown, though prior knowledge, if available, can be easily integrated into the model.

In a manner akin to [23], we model explicitly the sparsity of $\mathbf{A}$ by using the decomposition

$$\mathbf{A} = \mathbf{Z} \odot \mathbf{X} \tag{2}$$

where $\mathbf{Z} \in \{0, 1\}^{R\times N}$ is a matrix of Bernoulli variables that "labels" each material as *present* (active) or *absent* (inactive) in each pixel, $\mathbf{X} \in \boldsymbol{R}^{R\times N}$ is a matrix with positive entries that (jointly with $\mathbf{Z}$) quantifies the abundances, and $\odot$ denotes the Hadamard (term-wise) matrix product. This decomposition is particularly useful in Bayesian sparse regression problems because it allows eliciting separate statistical models for $\boldsymbol{a}_n$'s support (through modelling $\mathbf{Z}$), and for the values of the positive elements of $\boldsymbol{a}_n$ (through $\mathbf{X}$).

The next section presents a Bayesian model for estimating $\mathbf{Z}$ and $\mathbf{X}$ subject to the physical constraints discussed above. A key aspect of this model is that it will capture the fact that the pixels in which a material is active (or inactive) generally form spatial clusters (*i.e.*, exhibit spatial group sparsity). In difficult unmixing scenarios, exploiting this strong prior information can significantly improve estimation results, as will be shown in this paper.

III. BAYESIAN MODEL

This section presents the proposed the hierarchical Bayesian model for performing sparse source separation with collaborative supports. This model is defined by specifying the likelihood and the prior distribution of the parameters of interested $\mathbf{Z}$ and $\mathbf{X}$, as well as for the other unknown parameters in the model (e.g. $\boldsymbol{\sigma}^2$) that will be subsequently removed by marginalisation (*i.e.*, integrated out of the model's joint posterior distribution).

A. Likelihood

From the observation model (1) and the parametrisation of $\mathbf{A}$ described in (2), the likelihood of the image $\mathbf{Y}$ given the unknown parameters is

$$f(\mathbf{Y}|\mathbf{Z},\mathbf{X},\mathbf{M},\boldsymbol{\sigma}^2) = \prod_n f(\mathbf{y}_n|\boldsymbol{z}_n,\boldsymbol{x}_n,\mathbf{M},\boldsymbol{\sigma}^2) \quad (3)$$

$$= \prod_n p_{\mathcal{N}}(\mathbf{y}_n|\mathbf{M}(\boldsymbol{z}_n \odot \boldsymbol{x}_n),\boldsymbol{\Sigma}_0) \quad (4)$$

where $\boldsymbol{z}_n$ (resp. $\boldsymbol{x}_n$)is the nth column of $\mathbf{Z}$ (resp. $\mathbf{X}$) and $p_{\mathcal{N}}(\mathbf{y}_n|\mathbf{M}(\boldsymbol{z}_n \odot \boldsymbol{x}_n),\boldsymbol{\Sigma}_0)$ is the probability density function of a multivariate Gaussian vector with mean vector $\mathbf{M}(\boldsymbol{z}_n \odot \boldsymbol{x}_n)$ and diagonal covariance matrix $\boldsymbol{\Sigma}_0$.

B. Prior distribution of $\mathbf{Z}$

As explained previously, a key aspect of the proposed Bayesian model is to take into account the fact that the pixels in which a given material is present or absent typically form spatial groups or clusters (as opposed to being randomly distributed in space). From a modelling viewpoint, this can be represented by correlating the Bernoulli variables or labels $\boldsymbol{z}_{r,n}$ across the spatial dimension indexed by n. This can be achieved, for example, by stating that if a certain material is present (or absent) in a given pixel, this increases the probability of its presence (or absence) in neighbouring pixels. Taking into account that each material can exhibit its own spatial configuration, and the constraint that there must be at least one material present in each pixel, we propose to assign $\mathbf{Z}$ the following *truncated multivariate Ising Markov random field* prior

$$f(\mathbf{Z}|\boldsymbol{\beta}) = \frac{\psi(\mathbf{Z})}{C(\boldsymbol{\beta})} \exp\left[\sum_{r=1}^{R} \beta_r \phi_r(\mathbf{Z})\right] \quad (5)$$

with $\boldsymbol{\beta} = \{\beta_1, \ldots, \beta_R\}$,

$$\phi_r(\mathbf{Z}) = \sum_n \sum_{n' \in \mathcal{V}(n)} \delta(z_{r,n} - z_{r,n'}), \tag{6}$$

$$\psi(\mathbf{Z}) = \prod_n \max_r (z_{r,n}), \tag{7}$$

$$C(\boldsymbol{\beta}) = \sum_{\mathbf{Z}} \psi(\mathbf{Z}) \exp\left[\sum_{r=1}^{R} \beta_r \phi_r(\mathbf{Z})\right], \tag{8}$$

and where $\delta(\cdot)$ is the Kronecker delta function and $\mathcal{V}(n)$ denotes the set of neighbours of pixel (n) (in this paper we have used the 8-pixel neighbourhood). The hyper-parameters $\beta_1, \ldots, \beta_R$ act as regularization parameters that control the degree of spatial smoothness or regularity associated with each endmember, accounting for the fact that different materials may exhibit different spatial distributions.

To gain intuition about the proposed prior, we note that setting $\boldsymbol{\beta} = \mathbf{0}$ in (5) and considering the prior (9) for the matrix $\mathbf{X}$ leads to a Bernoulli-Gaussian type prior for the abundances which is closely related to the ℓ_0-ℓ_2 penalty often used for sparse regression. In the general scenarios where $\boldsymbol{\beta} \neq \mathbf{0}$, Eq. (5) introduces spatial correlations between the components of $\mathbf{Z}$ and, together with (9), leads to an ℓ_0-ℓ_2-type penalty promoting spatial group sparsity for the abundances. Prior (5) can also be understood as a collaborative prior, in the sense that, by assigning higher probabilities to configurations in which the supports of the abundance vectors $\boldsymbol{a}_n$ are spatially correlated, it pools or shares information between the sparse regressions that take place at each pixel. Crucially, by capturing the spatial correlations that occur naturally in hyperspectral images, this prior can improve significantly estimation results.

C. Prior distribution of $\mathbf{X}$

We assign the elements of $\mathbf{X}$ the following hierarchical prior distribution

$$x_{r,n}|s_r^2 \sim \mathcal{N}_{\mathbb{R}^+}\left(0, s_r^2\right) \tag{9}$$

$$s_r^2|\gamma, \nu \sim \mathcal{IG}(\gamma, \nu) \tag{10}$$

parametrised by some fixed hyper-parameters (γ, ν), and where we note that the prior on $x_{r,n}|s_r^2$ is truncated to $\mathbb{R}^+$ to reflect the positivity of $x_{r,n}$. This prior is very flexible and can be adjusted to represent a wide variety of prior beliefs (in all our experiments we used $\gamma = 2.1$ and $\nu = 1.1$, corresponding to a weakly informative prior for s_r^2 with 80% of its mass in $[0, 1]$.)

An important advantage of the hierarchical prior (9) is its natural capacity to encode prior dependences between the abundances. We expect the abundance coefficients associated with the same material to exhibit correlations, in particular in terms of their scale. This belief is encoded in (9) by defining a common parameter s_r^2 for each material or endmember $\mathbf{m}_r$, which is shared by all the abundances related to that material. Therefore, the hierarchical structure of (9) operates as a global pooling mechanism that shares information across the rows of $\mathbf{A}$ (i.e, the abundance coefficients associated to the rth material) to improve estimation accuracy.

Assuming an exchangeable structure where the abundances are prior independent given the hidden variables $\mathbf{s}^2 = [s_1^2, \ldots, s_R^2]^T$, we obtain the following following joint prior for $\mathbf{X}, \mathbf{s}^2$

$$f(\mathbf{X}, \mathbf{s}^2) = f(\mathbf{X}|\mathbf{s}^2) f(\mathbf{s}^2), \tag{11}$$

with $f(\mathbf{X}|\mathbf{s}^2) = \prod_{r,n} f(x_{r,n}|s_r^2)$ and $f(\mathbf{s}^2) = \prod_r f(s_r^2|\gamma, \nu)$. Also notice that by using the hierarchical structure (9) we obtain conjugate priors and hyper-priors for $x_{r,n}$ and s_r^2; this leads to inference algorithms with significantly better tractability and computational efficiency, which is crucial given the high dimensionality of $\mathbf{X}$.

Finally, notice that in the proposed model the spatial dependences in the abundance maps are encoded at the level of the abundance supports (though the prior on $\mathbf{Z}$), and not directly through the values of the abundances $\mathbf{A} = \mathbf{Z} \odot \mathbf{X}$ as it is the case in some previous models (see for example [19]). The motivation for this modelling choice is that modern high-resolution hyperspectral images often exhibit textures and fine detail that are not well described by models that promote smooth or piecewise-constant abundances maps. The experiments reported in this work show that the two approaches to modelling spatial correlation have complementary strengths and weaknesses, suggesting that to further improve estimation results future models should consider both levels of spatial dependences simultaneously.

D. Prior distribution of the noise variances $\boldsymbol{\sigma}^2$

In this paper we consider that there is no significant prior knowledge available about the values of the noise variances and assign each σ_ℓ^2 its non-informative Jeffreys prior [24]

$$f(\sigma_\ell^2) \propto \frac{1}{\sigma_\ell^2} \mathbf{1}_{\mathbb{R}^+}\left(\sigma_\ell^2\right), \quad \ell = 1, \ldots, L. \tag{12}$$

Note that in scenarios where prior knowledge about $\boldsymbol{\sigma}^2$ is available, this can be easily integrated into the model by replacing (12) with an inverse gamma conjugate priors with hyper-parameter values reflecting this prior knowledge.

E. Regularisation parameter β

A main advantage of Bayesian methods is that they allow estimating the appropriate amount of regularisation from data, thus freeing practitioners from the difficulty of setting regularisation parameters by cross-validation. Indeed, there are several Bayesian strategies for selecting the value of the regularisation parameter $\boldsymbol{\beta}$ in a fully automatic manner (see [25] for a recent detailed survey on this topic). In this paper we use the empirical Bayes technique recently proposed in [26], where the value of $\boldsymbol{\beta}$ is estimated by maximum marginal likelihood.

F. Joint posterior distribution of $\mathbf{Z}, \mathbf{X}, \boldsymbol{\sigma}^2, \mathrm{s}^2$

The structure of the proposed hierarchical Bayesian model is summarised in the directed acyclic graph (DAG) depicted below in Fig. 1.

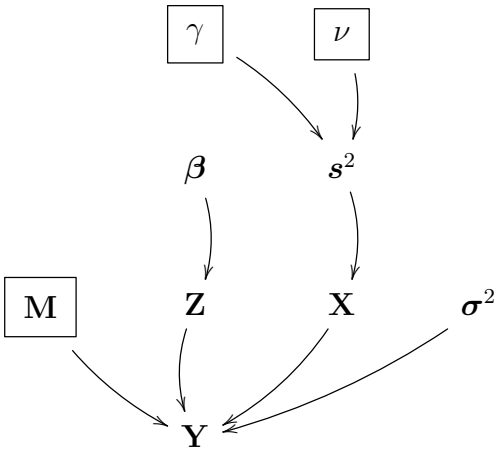

Fig. 1. Directed acyclic graph (DAG) of the proposed hierarchical Bayesian model (parameters with fixed values are represented using black boxes).

Using Bayes' theorem, and taking into account the conditional independences of the Bayesian model (see Fig. 1), the joint posterior distribution of $\mathbf{Z}, \mathbf{X}, \boldsymbol{\sigma}^2$ and s^2 given the observations $\mathbf{Y}$, the library of spectral signatures $\mathbf{M}$ and the model's fixed parameters $\boldsymbol{\beta}, \gamma$ and ν, is given by

$$f(\mathbf{Z}, \mathbf{X}, \boldsymbol{\sigma}^2, \mathrm{s}^2 | \mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu) \propto f(\mathbf{Y}|\mathbf{Z}, \mathbf{X}, \mathbf{M}, \boldsymbol{\sigma}^2) f(\mathbf{Z}|\boldsymbol{\beta}) f(\mathbf{X}|\mathrm{s}^2) f(\mathrm{s}^2|\gamma, \nu) f(\boldsymbol{\sigma}^2). \quad (13)$$

The following section presents a Monte Carlo algorithm to perform Bayesian inference using the proposed model.

IV. Bayesian inference using Gibbs sampling

The Bayesian model defined in Section III specifies a joint posterior distribution for the unknown parameters $\mathbf{Z}, \mathbf{X}, \boldsymbol{\sigma}^2, \mathrm{s}^2$ given the fixed quantities $\mathbf{Y}, \mathbf{M}, \gamma, \nu$ and the hyper-parameter $\boldsymbol{\beta}$ which is unknown but represented as a deterministic parameter (whose value will be tuned during the inference procedure). According to the Bayesian paradigm, this posterior distribution fully describes the information about the unknowns that is provided by the data and by the prior knowledge available. However, for unmixing applications it is necessary to summarise this posterior distribution in the form of point estimates; that is, to assign specific values for the unknown quantities of interest (in our problem the abundance vectors). Here we consider the following coupled Bayesian estimators that are particularly suitable for sparse regression problems: the marginal maximum *a posteriori* (MMAP) estimator [27], [28] for the support of the abundance vectors or "presence maps"

$$z_{r,n}^{MMAP} = \underset{z_{r,n}\in\{0,1\}}{\operatorname{argmax}} f(z_{r,n}|\mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu), \tag{14}$$

and, conditionally on the estimated supports, the minimum mean square error estimator of the abundances

$$a_{r,n}^{MMSE} = \mathrm{E}\left[x_{r,n}|z_{r,n} = \hat{z}_{r,n}^{MMAP}, \mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu\right], \tag{15}$$

where

$$f(z_{r,n}|\mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu) = \int f(\mathbf{Z}, \mathbf{X}, \boldsymbol{\sigma}^2, \mathrm{s}^2|\mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu)\mathrm{d}\mathbf{Z}_{\setminus z_{r,n}}\mathrm{d}\mathbf{X}, \mathrm{d}\boldsymbol{\sigma}^2\mathrm{ds},$$

with the matrix $\mathbf{Z}_{\setminus z_{r,n}}$ containing the remaining elements of $\mathbf{Z}$ once $z_{r,n}$ has been removed and where $\mathrm{E}[\cdot]$ denotes the expectation with respect to the conditional marginal density

$$f(x_{r,n}|z_{r,n}, \mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu) = \frac{\int f(\mathbf{Z}, \mathbf{X}, \boldsymbol{\sigma}^2, \mathrm{s}^2|\mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu)\mathrm{d}\mathbf{Z}_{\setminus z_{r,n}}\mathrm{d}\mathbf{X}_{\setminus z_{r,n}}, \mathrm{d}\boldsymbol{\sigma}^2\mathrm{ds}}{f(z_{r,n}|\mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu)}.$$

Note that the abundance estimator (15) is sparse by construction (*i.e.*, $\mathrm{E}\left[x_{r,n}|z_{r,n} = 0, \mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu\right] = 0$) and that, by marginalising out the other unknowns, it automatically takes into account the uncertainty about σ^2 and s^2.

The choice of specific Bayesian estimators to summarise the posterior distribution is a decision-theoretic problem that depends on the model and the application considered [27]. In the model described in this work, the two quantities of interest Z and X are very different in nature, and as a result their posterior information is best summarised with different estimators. The labels Z are binary variables that describe a quantitative aspect of the model; that is, they parametrise the hypotheses that materials are present or absent in each pixel. Selecting

the hypotheses with highest posterior probability leads to marginal MAP estimation [27], which in this case operates as a model-selection tool. Moreover, the conditional posterior distribution of X|Z represents the uncertainty regarding the abundance values for a specific model or configuration (in particular the one identified by marginal MAP estimation). To summarise this posterior distribution we use the MMSE estimator because it is optimal with respect to any quadratic loss function, and approximately optimal with respect to any convex loss function [27]. Furthermore, in cases where there is strong prior knowledge justifying the use of a sum-to-one constraint on the abundances, this information can be incorporated to the inferences by using an MMSE estimator constrained to the simplex. However, we have observed that this property does not hold on our images, in part because of the effects of mild shadows in the scene.

Computing (14) and (15) is challenging because it requires having access to the univariate marginal densities of $z_{r,n}$ and the joint marginal densities of $(x_{r,n}, z_{r,n})$, which in turn require computing the posterior (13) and integrating it over a very high-dimensional space. Fortunately these estimators can be efficiently approximated with arbitrarily large accuracy by Monte Carlo integration. Precisely, it is possible to compute (14) and (15) by first using a Markov Chain Monte Carlo (MCMC) computational method to generate samples asymptotically distributed according to (13), and subsequently using these samples to approximate the required marginal probabilities and expectations.

Here we propose a Gibbs sampler to simulate samples from (13), as this type of MCMC method is particularly suitable for models involving hidden Markov random fields [29, Chap. 10]. The output of this algorithm are two Markov chains of N_{MC} samples $\{\mathbf{X}^{(1)}, \ldots, \mathbf{X}^{(N_{\text{MC}})}\}$ and $\{\mathbf{Z}^{(1)}, \ldots, \mathbf{Z}^{(N_{\text{MC}})}\}$ that are asymptotically distributed according to the posterior distribution $f(\mathbf{Z}, \mathbf{X}|\mathbf{Y}, \mathbf{M}, \boldsymbol{\beta}, \gamma, \nu)$. The first N_{bi} samples of these chains correspond to the so-called *burn-in* transient period and should be discarded (the length of this period can be assessed visually from the chain plots or by computing convergence tests). The remaining $N_{\text{MC}} - N_{\text{bi}}$ of each chain are used to approximate the Bayesian estimators (14) and (15) as follows

$$\hat{z}_{r,n}^{MMAP} = \underset{u=\{0,1\}}{\operatorname{argmax}} \sum_{t=N_{\text{bi}}+1}^{N_{\text{MC}}} \delta\left(z_{r,n}^{(t)} - u\right), \tag{16}$$

and

$$\hat{a}_{r,n}^{MMSE} = \frac{\sum_{t=N_{\text{bi}}+1}^{N_{\text{MC}}} x_{r,n}^{(t)} \delta\left(z_{r,n}^{(t)} - \hat{z}^{MMAP}\right)}{\sum_{t=N_{\text{bi}}+1}^{N_{\text{MC}}} \delta\left(z_{r,n}^{(t)} - \hat{z}_{r,n}^{MMAP}\right)}. \tag{17}$$

Note that (16) and (17) converge to the true Bayesian estimators (14) and (15) as $N_{\text{MC}} \to \infty$. The remainder of this sections provides details about the main steps of the proposed Gibbs sampler, termed *Collaborative sparse Unmixing* (CSU) and summarised in Algo. 1 below. Note that for clarity the dependence of all distributions on the known fixed quantities $\mathbf{M}, \gamma, \nu$ and $\boldsymbol{\beta}$ is omitted.

ALGORITHM 1

Collaborative sparse Unmixing (CSU)

1: Fixed input parameters: $\mathbf{M}, K, \gamma, \nu$, number of burn-in iterations N_{bi}, total number of iterations N_{MC}
2: Initialization ($t = 0$)
- Set $\mathbf{X}^{(0)}, \mathbf{Z}^{(0)}, \boldsymbol{\sigma}^{2(0)}, \boldsymbol{s}^{2(0)}, \boldsymbol{\beta}^{(0)}$

3: Iterations ($1 \leq t \leq N_{\text{MC}}$)
4: Set $\mathbf{Z}^* = \mathbf{Z}^{(t-1)}$
5: **for** $n = 1 : N$ **do**
6: Sample $\mathbf{z}_n^* \sim f(\mathbf{z}_n^* | \mathbf{Y}, \mathbf{Z}^*_{\setminus \mathbf{z}_n^*} \mathbf{X}^{(t-1)}, \boldsymbol{\sigma}^{2(t-1)}, \boldsymbol{s}^{2(t-1)})$ in (18)
7: **end for**
8: Set $\mathbf{Z}^{(t)} = \mathbf{Z}^*$
9: Sample $\mathbf{X}^{(t)} \sim f(\mathbf{X} | \mathbf{Y}, \mathbf{Z}^{(t)}, \boldsymbol{\sigma}^{2(t-1)}, \boldsymbol{s}^{2(t-1)},)$ in (20)
10: Sample $\boldsymbol{\sigma}^{2(t)} \sim f(\boldsymbol{\sigma}^2 | \mathbf{Y}, \mathbf{X}^{(t)}, \mathbf{Z}^{(t)}, \boldsymbol{s}^{2(t-1)},)$ in (22)
11: Sample $\boldsymbol{s}^{2(t)} \sim f(\boldsymbol{s}^2 | \mathbf{Y}, \mathbf{X}^{(t)}, \mathbf{Z}^{(t)}, \boldsymbol{\sigma}^{2(t)})$ in (23)
12: Update $\boldsymbol{\beta}^{(t)} \leftarrow \boldsymbol{\beta}^{(t-1)}$ using [26].
13: Set $t = t + 1$.

A. *Sampling the label matrix* $\mathbf{Z}$

The label matrix $\mathbf{Z}$ is updated pixel-wise by iteratively simulating from the distribution of the labels at each pixel given the other pixels. Precisely, the distribution of the vector $\mathbf{z}_n$ given the matrix $\mathbf{Z}_{\setminus \mathbf{z}_n}$ containing the remaining elements of $\mathbf{Z}$ once $\mathbf{z}_n$ is removed is given by

$$\begin{aligned} P(\mathbf{z}_n | \mathbf{Y}, \mathbf{Z}_{\setminus \mathbf{z}_n}, \mathbf{X}, \boldsymbol{\sigma}^2, \mathrm{s}^2) &\propto f(\mathbf{y}_n | \mathbf{z}_n, \mathbf{x}_n, \boldsymbol{\sigma}^2) f(\mathbf{z}_n | \mathbf{Z}_{\setminus \mathbf{z}_n}) \qquad (18) \\ &\propto \exp\left[2 \sum_{r=1}^{R} \beta_r \sum_{n' \in \mathcal{V}(n)} \delta(z_{r,n} - z_{r,n'}) \right] \exp\left[-\frac{1}{2\sigma^2} \tilde{\mathbf{y}}_n^T \boldsymbol{\Sigma}_0^{-1} \tilde{\mathbf{y}}_n \right] \qquad (19) \end{aligned}$$

where $\tilde{\mathbf{y}}_n = \mathbf{y}_n - \mathbf{M}(\mathbf{x}_n \odot \mathbf{z}_n)$, for $||\mathbf{z}_n||_0 > 0$ and $P(\mathbf{z}_n | \mathbf{Y}, \mathbf{Z}_{\setminus \mathbf{z}_n}, \mathbf{X}, \boldsymbol{\sigma}^2, \mathrm{s}^2) = 0$ otherwise. Algorithmically, this simulation step can be achieved by indexing the $2^R - 1$ admissible

configurations of $\mathbf{z}_n$ and then randomly selecting a specific one with probability defined in (18).

B. Sampling $\mathbf{X}$

The conditional distribution of $\mathbf{X}$ given the other unknown parameters can be factorised pixel-wise as a product of N marginal distributions

$$f(\mathbf{X}|\mathbf{Y},\mathbf{Z},\boldsymbol{\sigma}^2,\boldsymbol{s}^2)=\prod_{n=1}^{N} f(\mathbf{x}_n|\mathbf{y}_n,\mathbf{z}_n,\boldsymbol{\sigma}^2,\boldsymbol{s}^2), \tag{20}$$

that can be efficiently sampled independently and in parallel

$$\mathbf{x}_n|\mathbf{y}_n,\mathbf{z}_n,\boldsymbol{\sigma}^2,\boldsymbol{s}^2 \sim \mathcal{N}_{\Omega}(\bar{\mathbf{x}}_n,\boldsymbol{\Sigma}_n), \tag{21}$$

where $\Omega=(\mathbb{R}^+)^R$ is the positive orthant of $\mathbb{R}^R$ and

$$\begin{aligned}\boldsymbol{\Sigma}_n &= \left(\mathbf{D}_n\mathbf{M}^T\boldsymbol{\Sigma}_0^{-1}\mathbf{M}\mathbf{D}_n+\mathbf{S}^{-1}\right)^{-1}\\ \bar{\mathbf{x}}_n &= \boldsymbol{\Sigma}_n\mathbf{D}_n\mathbf{M}^T\boldsymbol{\Sigma}_0^{-1}\mathbf{y}_n,\end{aligned}$$

and where $\mathbf{S}=\text{diag}(\mathbf{s}^2)$ and $\mathbf{D}_n=\text{diag}(\mathbf{z}_n)$ are diagonal matrices with diagonal elements given by $\mathbf{s}^2$ and $\mathbf{z}_n$. For completeness, the derivation of 21 is provided in Appendix. In this paper we use the method [30] to simulate efficiently from (21).

C. Sampling the noise variances $\boldsymbol{\sigma}^2$

It can be easily shown that the noise variances are (conditioned on the other parameters) a posteriori dependent and can thus be updated in a parallel manner. Precisely, the conditional distribution associated with σ_ℓ^2 has a simple closed form expression and is given by

$$\sigma_\ell^2|\mathbf{Y},\mathbf{X},\mathbf{Z},\boldsymbol{s}^2 \sim \mathcal{IG}\left(N/2, E_{\sigma_\ell}\right), \tag{22}$$

with $E_{\sigma_\ell}=\|\tilde{\mathbf{y}}_{\ell,:}\|_2^2/2$ ($\tilde{\mathbf{y}}_{\ell,:}$ denotes the ℓth row of the $L\times N$ matrix $\widetilde{\mathbf{Y}}=[\tilde{\mathbf{y}}_1,\ldots,\tilde{\mathbf{y}}_N]$).

D. Sampling the hyperparameter vector $\mathbf{s}^2$

The conditional distribution of $\mathbf{s}^2$ can be factorised endmember-wise as a product of R independent marginal distributions

$$f(\boldsymbol{s}^2|\mathbf{Y},\mathbf{X},\mathbf{Z},\boldsymbol{\sigma}^2)=\prod_{r=1}^{R} f(s_r^2|\mathbf{X}) \tag{23}$$

that can be easily sampled independently and in a parallel manner

$$s_r^2|\mathbf{X} \sim \mathcal{IG}\left(\frac{N}{2}+\gamma, \frac{\sum_n \mathbf{x}_{r,n}^2}{2}+\nu\right). \tag{24}$$

V. Validation with synthetic data

This section demonstrates the proposed methodology on a series of experiments conducted using synthetic data. An applications to a real hyperspectral image is reported in Section VI.

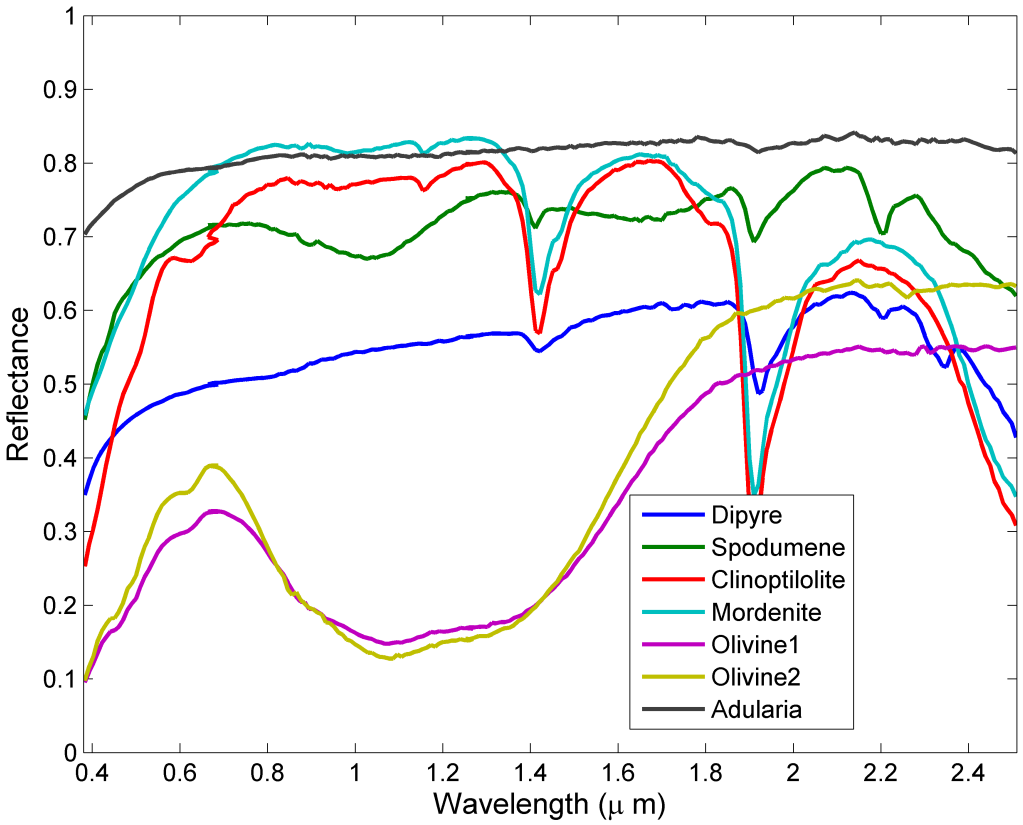

Fig. 2. Seven endmembers from the USGS spectral library.

A. Data sets

The performance of the proposed collaborative sparse unmixing (CSU) method is first evaluated on two synthetic images I_1 and I_2 of size 100×100 pixels and $L = 224$ spectral bands. There are $R_0 = 5$ endmembers present in the images which correspond to the minerals Dipyre, Spodumene, Clinoptilolite, Mordenite and Olivine 1. Their spectral signatures have been obtained from the USGS spectral library [31] and are depicted in Fig. 2. Note that the angles between the spectral signatures are between $3.01°$ and $3.05°$ (*i.e.*, the endmembers are highly correlated with mutual coherence [32] that equals to $M = 0.9986$), making the unmixing problem very challenging. The support maps that determine the spatial distribution

of each material have been generated by simulating from the prior model (5) with $\boldsymbol{\beta} = [0.2; 0.275; 0.35; 0.425; 0.5]^T$, and are depicted in the top row of Fig. 3. For both images the matrix $\mathbf{X}$ has been generated by sampling from (9) with $s_r = 0.3, \forall r$. Finally, the average signal-to-noise ratio (SNR) for the images I_1 and I_2 are approximately 30dB ($\sigma_\ell^2 = 8.10^{-4}, \forall \ell$) and 20dB ($\sigma_\ell^2 = 8.10^{-3}, \forall \ell$), respectively.

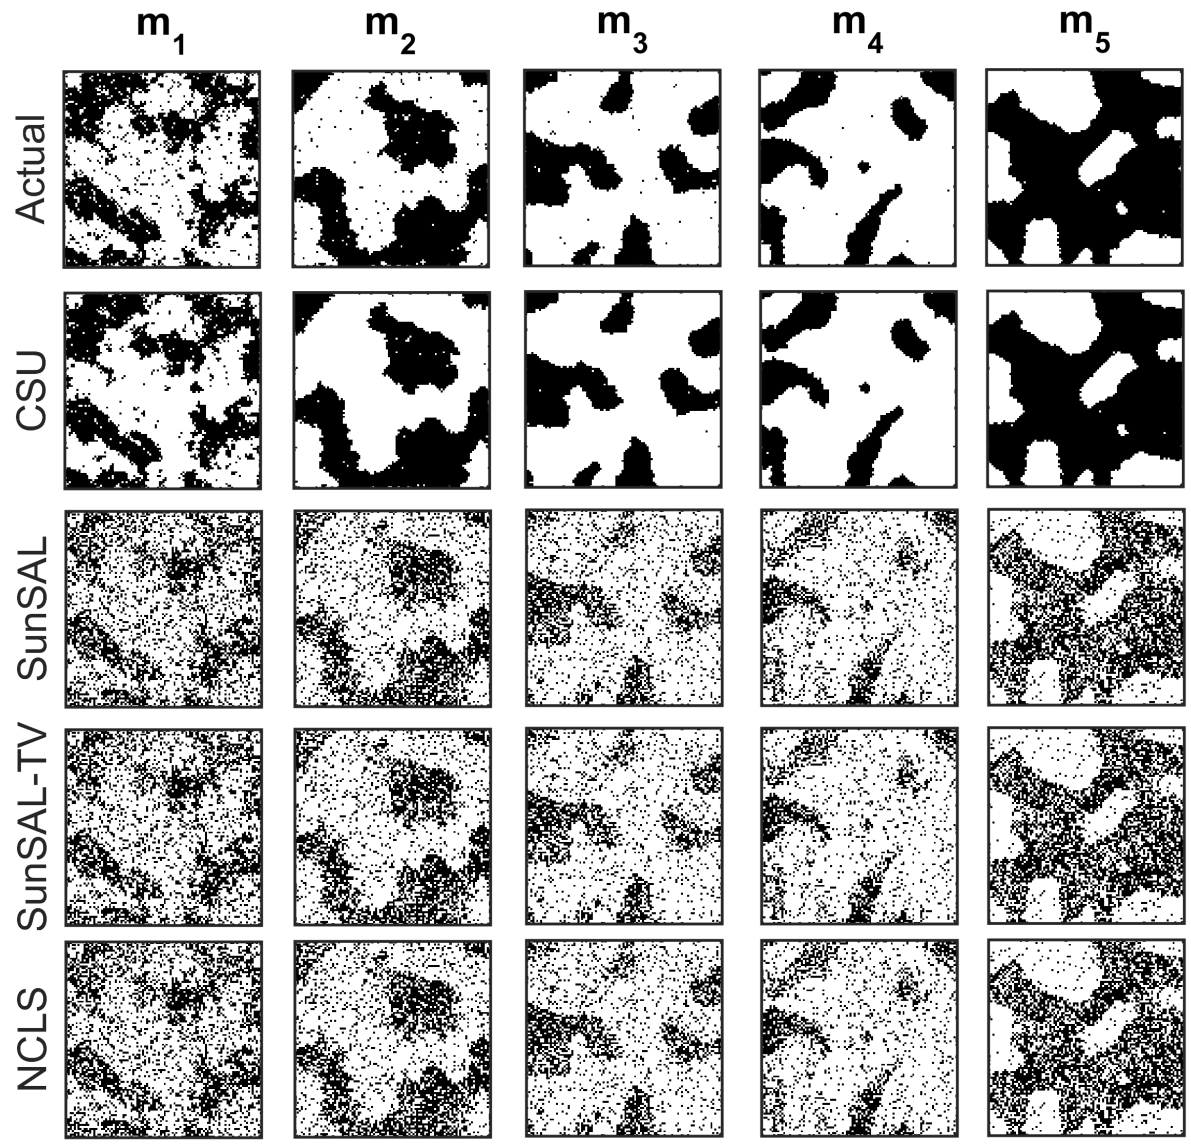

Fig. 3. Top row: Active support maps of the $R_0 = 5$ endmembers associated with I_1 and I_2. Active support maps estimated with the proposed CSU algorithm, SunSAL, SunSAL-TV, and NCLS ($\rho = 0.01$) for I_2 (top to bottom). White (resp. black) pixels correspond to regions where a component is present (resp. absent).

B. Supervised unmixing

In this first experiment we consider that the materials present in the images I_1 and I_2 are perfectly known and we estimate the abundance vectors with CSU, CLSunSAL and NCLS. The CSU algorithm has been implemented using $N_{\text{MC}} = 3000$ and $N_{\text{bi}} = 1000$, and by allowing the algorithm to self-adjust the regularisation parameter $\boldsymbol{\beta}$ with the technique proposed in [26]. The estimated values of $\boldsymbol{\beta}$ are $\boldsymbol{\beta} = [0.19, 0.28, 0.33, 0.37, 0.44]^T$ and $\boldsymbol{\beta} = [0.20, 0.28, 0.33, 0.33, 0.44]^T$ for I_1 and I_2, respectively. For comparison we use the state-of-the-art sparse regression algorithm SunSAL [10] and the SunSAL algorithm using the total variation regularization (SunSAL-TV) [19], whose regularisation parameter values are adjusted to provide the best abundance estimates (in terms of RMSE) for each scenario. For completeness we also compare with the widely used Non-negatively Constrained Least-Squares algorithm (NCLS) [5], which solves a maximum likelihood problem and does not exploit any prior information about the abundance vectors.

The associated computation times for a Matlab implementation on a 3.0GHz Intel Xeon quad-core workstation are provided in Table I. The estimated supports obtained with CSU, SunSAL, SunSAL-TV and NCLS for I_2 (lowest SNR) are depicted in Fig. 3. The presence maps associated with I_1 are similar and are not presented here due to space constraints. For SunSAL, SunSAL-TV and NCLS, the detection maps have been obtained by thresholding the estimated abundances with a threshold arbitrarily set to $\rho = 0.01$. We observe that the results obtained with CSU are in good agreement with the ground truths. On the other hand, the abundances obtained with SunSAL, SunSAL-TV and NCLS are significantly less accurate, thus confirming that taking into account the spatially correlations between the supports abundance vectors is key to achieving accurate estimation results in this scenario. This valuable prior knowledge is all the more important in low SNR scenarios such as the one depicted in Fig. 3.

For numerical comparison, we computed the root mean square error (RMSE) $\text{RMSE}_n = \sqrt{\|\hat{\boldsymbol{a}}_n - \boldsymbol{a}_n\|^2}$, that quantifies the average accuracy of the estimated abundances $\hat{\boldsymbol{a}}_n$ with respect to the truth $\boldsymbol{a}_n$ at the n-th pixel. We also consider the abundance angle distance (AAD) $\text{AAD}_n = \cos^{-1}\left(\frac{\hat{\boldsymbol{a}}_n^T \boldsymbol{a}_n}{\|\hat{\boldsymbol{a}}_n\|_2 \|\boldsymbol{a}_n\|_2}\right)$, which is not sensitive to scaling factors between actual and estimated abundance vectors. The lower the RMSEs and AADs, the better the abundance estimation performance.

The first five rows of Table II show the average RMSEs and AADs, for CSU, SunSAL,

SunSAL-TV, NCLS, and the oracle NCLS (o-NCLS), which consists of applying NCLS using only the active materials in each pixel (*i.e.*, with perfect knowledge of the support of the abundance vectors). We observe that CSU provides significantly more accurate estimations than SunSAL, SunSAL-TV and NCLS, achieving average RMSEs and AADs that are close to the oracle. Again, the good performance of CSU results from the collaboration between pixels introduced by the prior model (5). By comparing the first and second columns of Table II we confirm that this prior knowledge becomes all the more important as the noise level increases. These results show that using only sparsity (SunSAL) does not lead to significant improvements, perhaps because the ℓ_1 regularisation is not appropriate for our images, or because it would require adapting the regularisation parameters to each endmember (this is done automatically in CSU). We also observe that SunSAL-TV does not achieve significantly better results, perhaps due to the regions where the abundances exhibit fluctuations that are not well modelled by the TV. As mentioned previously, when both the abundance supports and the abundance values exhibit significant spatial dependencies, one should consider methods that account for this prior information (e.g., the SunSAL-TV method if the abundance maps are expected to be piecewise constant across the image). However, in many high-resolution images and regions in images, only the abundance supports exhibit significant spatial dependencies; in this case regularising the abundance values may lead to over-smooth estimation results. In this paper we show that by modelling these dependencies it is possible to improve unmixing performance with respect to methods that assume that abundances are uncorrelated.

Finally, in order to further highlight the good performance of the proposed prior model we have also computed the reconstruction error (RE) defined as $\text{RE}_n = \sqrt{\|\mathbf{M}\hat{\boldsymbol{a}}_n - \mathbf{y}_n\|^2}$. Note that this error essentially measures the likelihood of the abundance estimates given the observed image $\mathbf{Y}$ and assuming the noise is identically distributed in the L spectral bands. However, because the SU problem is not well-posed, the capacity of the likelihood to identify good solutions is severely limited and the additional information provided by the prior model is key to deliver accurate estimation results. In these scenarios Bayesian methods, which combine observed and prior information, can greatly outperform other estimation techniques. The average RE for each method is reported in Table III. We observe that all the algorithms considered exhibit very similar REs. However, we know from Table II that CSU outperforms significantly the other methods in terms of abundance estimation accuracy. The contrast between these two figures of merit confirms that the superior performance of CSU is directly

TABLE I

COMPUTATIONAL TIME (IN SECONDS): SYNTHETIC IMAGES .

		I_1	I_2
o-NCLS		5.08	14.54
R=5	CSU	3360	3420
	SunSAL	0.62	0.57
	SunSAL-TV	56	55
	NCLS	0.67	0.63
R=7	CSU	51330	52500
	SunSAL	0.55	0.60
	CLSunSAL	0.62	0.57
	SunSAL-TV	56	53
	NCLS	2.03	3.40

TABLE II

AVERAGE RMSES AND AADS: SYNTHETIC IMAGES .

		Av. RMSEs ($\times 10^{-2}$)		Av. AADs ($\times 10^{-2}$)	
		I_1	I_2	I_1	I_2
o-NCLS		6.01	17.20	7.56	20.91
R=5	CSU	6.30	17.05	8.07	21.32
	SunSAL	8.50	23.52	12.32	32.35
	SunSAL-TV	8.26	22.96	11.57	31.55
	NCLS	8.50	23.54	11.57	31.52
R=7	CSU	6.61	17.36	8.44	21.69
	SunSAL	13.13	31.87	19.04	43.87
	CLSunSAL	13.11	31.85	18.73	43.87
	SunSAL-TV	13.06	31.68	18.01	43.45
	NCLS	13.13	31.89	18.01	43.86

related to the proposed prior model, which captures the correlations between the supports of the abundance vectors and effectively introduces a means of collaboration that allows sharing information between pixels and increasing robustness to noise.

TABLE III

AVERAGE REs ($\times 10^{-2}$): SYNTHETIC IMAGES .

		I_1	I_2
R=5	CSU	2.81	8.89
	SunSAL	2.80	8.86
	SunSAL-TV	2.80	8.86
	NCLS	2.80	8.86
R=7	CSU	2.81	8.89
	SunSAL	2.80	8.87
	CLSunSAL	2.80	8.87
	SunSAL-TV	2.80	8.86
	NCLS	2.80	8.86

C. Semi-supervised unmixing

We now consider that two additional endmembers are incorrectly included in the library $\mathbf{M}$, although they are not present in the scene. To make the SU problem particularly challenging, the two additional endmembers are Olivine 2 and Adularia, whose spectral signatures are highly correlated with the signatures of the other endmembers. For example, note from Fig. 2 that discriminating between Olivine 1 and Olivine 2 is very difficult. We contrast CSU with SunSAL, SunSAL-TV and NCLS (CSU is implemented using $N_{\text{MC}} = 7000$ and $N_{\text{bi}} = 5000$ and the estimated values of $\boldsymbol{\beta}$ are $\boldsymbol{\beta} = [0.20, 0.28, 0.35, 0.40, 0.45, 0.47, 0.47]^T$ and $\boldsymbol{\beta} = [0.20, 0.29, 0.36, 0.37, 0.45, 0.46, 0.42]^T$ for I_1 and I_2, respectively.). We also consider CLSunSAL [11] which promotes group-sparsity for the estimated abundances.

Fig. 4 shows the support estimations obtained with CSU, SunSAL, CLSunSAL and SunSAL-TV for I_2 (lowest SNR) with the $R = 7$ endmembers (two of which are absent from the scene). The results obtained for I_1 follow the same trend and are not presented here. Again, we observe that the results obtained with CSU are in very good agreement with the ground truths, whereas the results obtained with the other methods are significantly less accurate, even when tuning the regularization parameter. More importantly, we observe that CSU has successfully detected that two endmembers (Olivine 2 and Adularia) are not present in the scene, in spite of the strong similarities between Olivine 1 and Olivine 2.

Finally, the five bottom rows of Table II show the RMSEs and AADs obtained with CSU, SunSAL, SunSAL-TV, CLSunSAL and NCLS for I_1 and I_2 with $R = 7$. We observe that

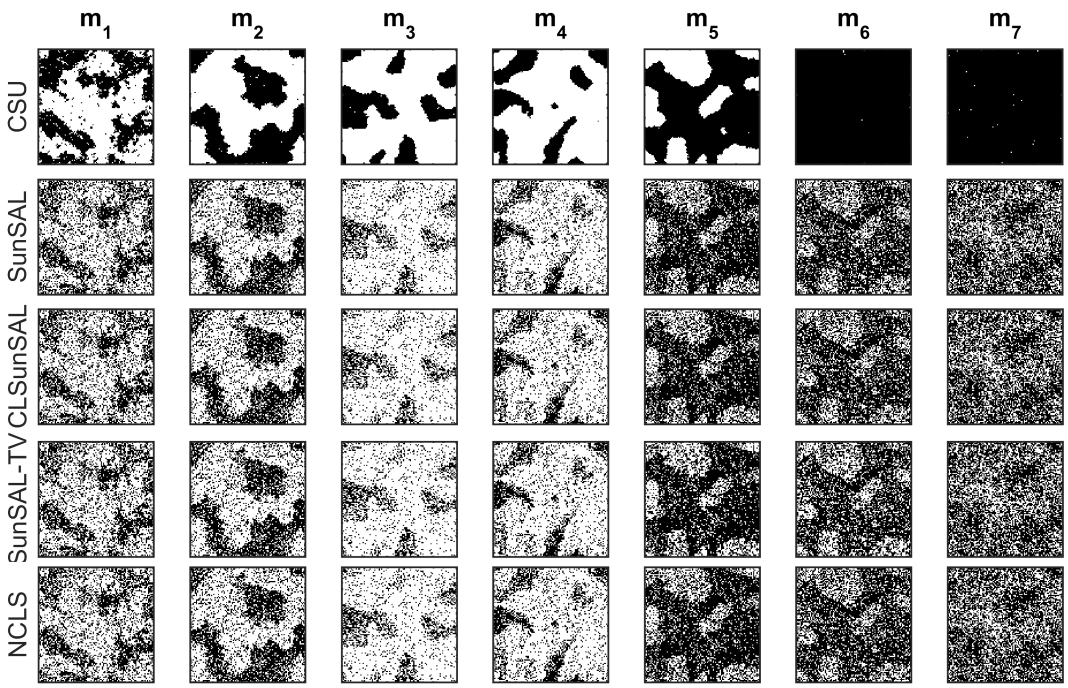

Fig. 4. Active support maps estimated with the proposed CSU algorithm, SunSAL, CLSunSAL, SunSAL-TV and NCLS (top to bottom, $\rho = 0.01$) for I_2 and $R = 7$ endmembers. White (resp. black) pixels correspond to regions where a component is present (resp. absent).

CSU is robust to the presence of the two redundant endmembers and is able to accurately discriminate between Olivine 1 and Olivine 2. On the other hand, SunSAL, SunSAL-TV, CLSunSAL and NCLS have difficulties detecting the true supports of the abundance vectors and produce estimation results that are significantly less accurate, in particular in low SNR conditions. Again, the superior performance of CSU is directly related to the prior model (5) that captures the spatial correlations of the mixture support in the image and regularises the supports of the abundance vectors. Notice that the group-sparsity model operating in CLSunSAL is unable to identify correctly the endmembers present in the images (due to the high correlation between the two Olivine spectra). Increasing the CLSunSAL regularisation parameter leads to a correct identification of the endmembers, but at the expense of a severe degradation in estimation performance. Similarly, increasing the SunSAL-TV regularisation parameter improves the detection of absent endmembers, but degrades overall estimation performance. Notice that CSU does not suffer from these drawbacks because it self-adjusts the scale parameters ($\mathbf{s}^2$) and spatial correlation parameters ($\boldsymbol{\beta}$) automatically during the inference procedure. For completeness, Fig. 5 presents the abundances of the endmembers

$5, 6$ and 7 (which are the hardest to discriminate) estimated with the different algorithms for I_2.

VI. Application to real a hyperspectral image

This section presents an application of the proposed CSU method to a real hyperspectral image acquired by the Hyspex hyperspectral scanner over Villelongue, France ($00°03'$W and $42°57'$N). This images was acquired in 2010 as part of the Madonna project, and is composed of $L = 160$ spectral bands covering from the visible to near infrared spectrum and with a spatial resolution of 0.5m (for more details about the data acquisition and pre-processing steps see [33].). This dataset has previously been studied in [22], [33], [34] and is mainly composed of forested and urban areas. Due to the high spatial resolution of the image, materials present in a pixel are likely to be also present in its neighbour pixels; that is, we expect significant spatial correlations between the mixture supports. Here we evaluate the proposed unmixing method on the region of interest of size 100×100 pixels depicted in Fig. 6. This region is composed mainly of trees and grass, with $R = 6$ endmembers related to soil, two types of grass, two types of trees, and an additional endmember modelling attenuation effects mainly related to shade. The spectral signatures for these endmembers have been extracted manually from the data by using our prior knowledge about the scene [33]. Note that for this image and within the LMM framework, the attenuation effects have to be modelled as an additional endmember because they operate differently on the different spectral bands of the reflectance spectra (otherwise, if the attenuations acted as a scaling factor affecting all the bands similarly, we would be able to handle this by relaxing the abundance sum-to-one constraint).

Fig. 7 shows the presence maps for each material estimated with CSU (using $N_{\text{MC}} = 5000$, $N_{\text{bi}} = 1000$ and by allowing the algorithm to self-adjust the regularisation parameter β with the technique [26]) and with NCLS, SunSAL and SunSAL-TV (whose parameters have been optimised in a fully supervised manner to obtain the sparser and/or smoother abundance maps without degrading significantly the reconstruction error, and by using a detection threshold of $\rho = 0.01$). Fig. 8 shows the total number of materials present in each pixels, as estimated by each method. Lastly, Fig. 10 shows the abundance maps estimated with each method. These figures clearly show that CSU provides sparser and spatially smoother presence (support) maps than the other methods. We also observe that, as expected, SunSAL-TV provides smoother abundance maps than the other methods (see Fig. 10). This property of

SunSAL-TV is beneficial in regions where the materials are absent; however, where materials are present, it may smooth out fine detail (e.g., textures). Since no abundance ground truth is available for this data set, it is difficult to determine which algorithm provides the more accurate results. However, due to the high spatial resolution of the image and the structure of the materials considered, it is reasonable to assume that abundances will exhibit some degree of local heterogeneity, and that CSU and SunSAL-TV will outperform each other in different ways. Moreover, from Figs. 7 and Fig. 10 we also note that CSU detects a higher lever of attenuation effects than SunSAL and SunSAL-TV, particularly in pixels containing vegetation where strong shadowing effects occur.

Finally, as a model checking analysis, we plot in Fig. 9 the marginal noise variances estimated by CSU (i.e., diagonal elements of the full estimated covariance matrix). For comparison we also include the estimates obtained with Hysime [35]. We observe that the estimation results produced by CSU are very smooth and physically realistic, suggesting good fit to data (recall that CSU assumes that the noise variances are prior independent, hence the smoothness of these posterior estimates arises from the data and from correlations with other model parameters). We also note that, although the estimation results produced by CSU and Hysime are different, the noise variances follow similar profiles. Lastly, for completeness Table IV reports the reconstruction errors associated with CSU, NCLS, SunSAL and SunSAL-TV. We observe that all methods achieve comparable reconstruction errors, with NCLS and SunSAL having a slightly lower reconstruction error than CSU and SunSAL-TV. This difference is likely due to the fact that NCLS and SunSAL do not enforce spatial regularity, and are hence more prone to overfitting the data.

TABLE IV

AVERAGE RECONSTRUCTION ERRORS (ARE): REAL IMAGE.

Unmixing algo.	ARE ($\times 10^{-2}$)
CSU	7.14
SunSAL	6.98
SunSAL-TV	7.00
NCLS	6.88

VII. Conclusion

This paper presented a new Bayesian method for linear unmixing of hyperspectral image that is based on a collaborative sparse regression formulation. The main novelty is a Bayesian model for sparse regression that takes into account the fact that the supports of the abundance vectors are spatially correlated. This prior information is encoded in the model by using a truncated multivariate Ising Markov random field model, which also accounts for the facts that pixels cannot be empty (*i.e.*, there is at least one material per pixel) and that each material in the scene may require a different amount of spatial regularisation. The proposed Bayesian model also takes into consideration that material abundances are non-negative quantities and that the level of noise contaminating the image may be unknown. Following on from this, we presented a Markov chain Monte Carlo algorithms to perform Bayesian inference with this model and compute the statistical estimators of interest. Precisely, we proposed a Gibbs sampler that allows estimating the probabilities that materials are present or absence in each pixel, and, conditionally on any given configuration (typically the maximum *a posteriori*), computing the MMSE estimates of the abundance vectors. A remarkable characteristic of the proposed inference algorithm is that it self-adjusts the amount of regularity enforced by the random field, thus relieving practitioners from the difficult task of setting regularisation parameters by cross-validations. Finally, the good performance of the proposed methodology was demonstrated through a series of experiments with synthetic and real data and comparisons with other algorithms from the literature.

Due to its computational complexity, the algorithm presented in this paper can only be directly applied to problems with small numbers of endmembers (*e.g.*, $R \leq 25$). For problems with larger libraries it is computationally more efficient to generate samples from a relaxed posterior in which the "non-empty pixel" constraint (7) of the MRF is removed, and then reintroduce this constraint by importance sampling. Another possibility for problems with large libraries is to use the MUSIC-CSR algorithm [12] as a pre-processing step to identify endmembers that are absent from the scene, and then apply our method using a pruned library.

As mentioned previously, previous works have considered that the values of the abundance vectors are spatially correlated, whereas we considered correlations between their supports (*i.e.*, material presence and absence patterns). Investigating new models that exploit both types of correlations is an important perspective for future work. For many applications or

hyperspectral datasets, it makes sense to consider additional abundance constraints, such as the sum-to-one constraint to further improve unmixing results. Embedding this constraint within ℓ_0-type sparse regression models is a challenging problem that is currently under investigation. Another perspective for future work is to investigate more sophisticated spatial models that describe hyperspectral images more accurately, in particular in complex scenes with numerous materials and non-linear effects. Finally, we believe that the methodology presented in this paper could be interesting for other regression problems that exhibit structured sparsity and intend to further investigate this is the future.

ACKNOWLEDGMENTS

The authors would like to thank Prof. Jean-Yves Tourneret and Dr Nicolas Dobigeon, from the University of Toulouse, IRIT-ENSEEIHT, France, for interesting discussion regarding this work.

APPENDIX: ON THE CONDITIONAL DISTRIBUTION OF THE ABUNDANCE VALUES

Due to the conjugacy of the priors (9) and the likelihood (3), the posterior distribution

$$f(\mathbf{x}_n|\mathbf{y}_n, \mathbf{z}_n, \boldsymbol{\sigma}^2, \boldsymbol{s}^2) \propto f(\mathbf{y}_n|\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\sigma}^2) f(\mathbf{x}_n|\boldsymbol{s}^2)$$

is a multivariate Gaussian distribution restricted to Ω, i.e., the positive orthant of $\mathbb{R}^R$. Moreover,$\forall \mathbf{x}_n \in \Omega$

$$\begin{aligned}
&\log\left(f(\mathbf{x}_n|\mathbf{y}_n, \mathbf{z}_n, \sigma^2, \boldsymbol{s}^2)\right) \\
&\quad = \log\left(f(\mathbf{y}_n|\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\sigma}^2)\right) \log\left(f(\mathbf{x}_n|\boldsymbol{s}^2)\right), \\
&\quad = -\frac{\tilde{\mathbf{y}}_n^T \boldsymbol{\Sigma}_0^{-1} \tilde{\mathbf{y}}_n}{2} - \frac{\mathbf{x}_n^T \mathbf{S}^{-1} \mathbf{x}_n}{2} + c_1 \\
&\quad = -\frac{1}{2}\mathbf{x}_n^T \left(\mathbf{S}^{-1} + \mathbf{D}_n \mathbf{M}^T \boldsymbol{\Sigma}_0^{-1} \mathbf{M} \mathbf{D}_n\right) \mathbf{x}_n \\
&\qquad + \mathbf{x}_n^T \mathbf{D}_n \mathbf{M}^T \boldsymbol{\Sigma}_0^{-1} \mathbf{y}_n + c_2
\end{aligned}$$

where $\mathbf{S} = \text{diag}(\mathbf{s}^2)$ and $\mathbf{D}_n = \text{diag}(\mathbf{z}_n)$ are diagonal matrices with diagonal elements given by $\mathbf{s}^2$ and $\mathbf{z}_n$ and c_1, c_2 are real constants (independent of $\mathbf{x}_n$). By identification, we obtain $\boldsymbol{\Sigma}_n^{-1} = \mathbf{D}_n \mathbf{M}^T \boldsymbol{\Sigma}_0^{-1} \mathbf{M} \mathbf{D}_n + \mathbf{S}^{-1}$, $\boldsymbol{\Sigma}_n^{-1} \bar{\mathbf{x}}_n = \mathbf{D}_n \mathbf{M}^T \boldsymbol{\Sigma}_0^{-1} \mathbf{y}_n$ and finally

$$\begin{aligned}
\boldsymbol{\Sigma}_n &= \left(\mathbf{D}_n \mathbf{M}^T \boldsymbol{\Sigma}_0^{-1} \mathbf{M} \mathbf{D}_n + \mathbf{S}^{-1}\right)^{-1} \\
\bar{\mathbf{x}}_n &= \boldsymbol{\Sigma}_n \mathbf{D}_n \mathbf{M}^T \boldsymbol{\Sigma}_0^{-1} \mathbf{y}_n,
\end{aligned}$$

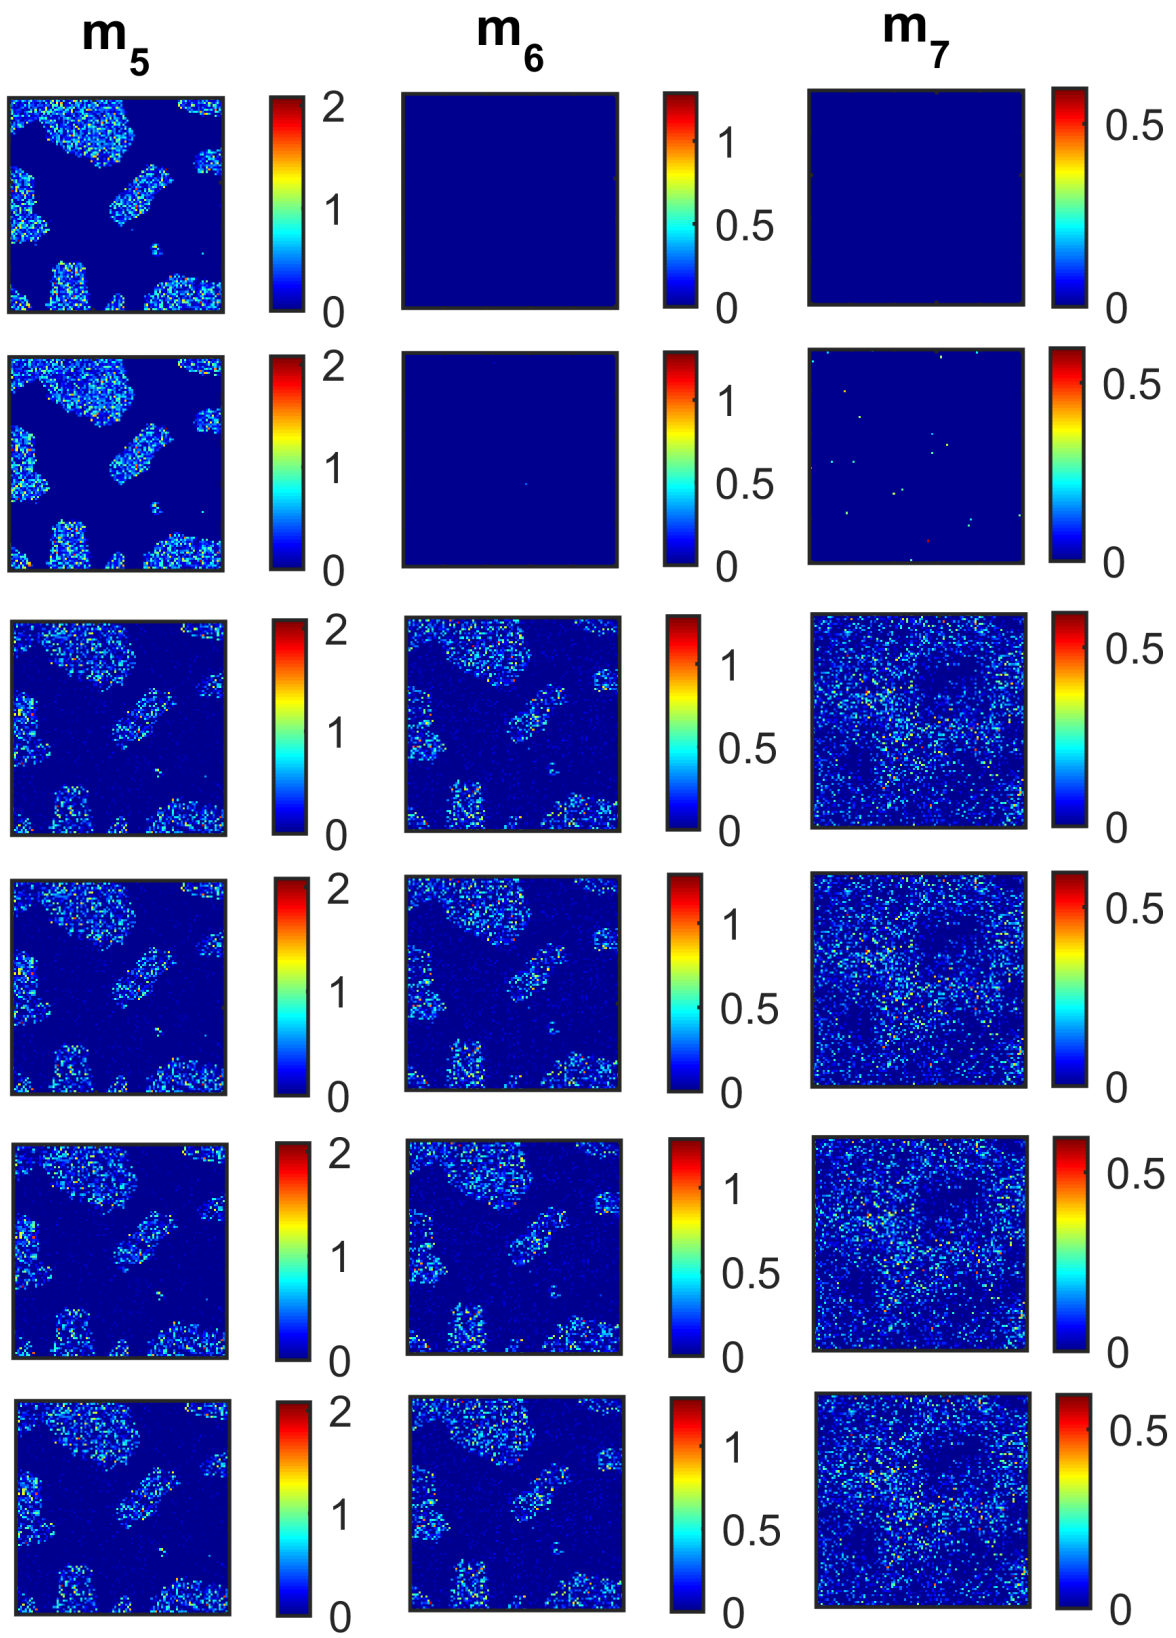

Fig. 5. Top: Actual abundance maps of the endmembers $5, 6$ and 7 for I_2. Abundance maps estimated by CSU, SunSAL, CLSunSAL, SunSAL-TV and NCLS (top to bottom).

Fig. 6. True color image of the Villelongue area (left) and sub-image of interest (right).

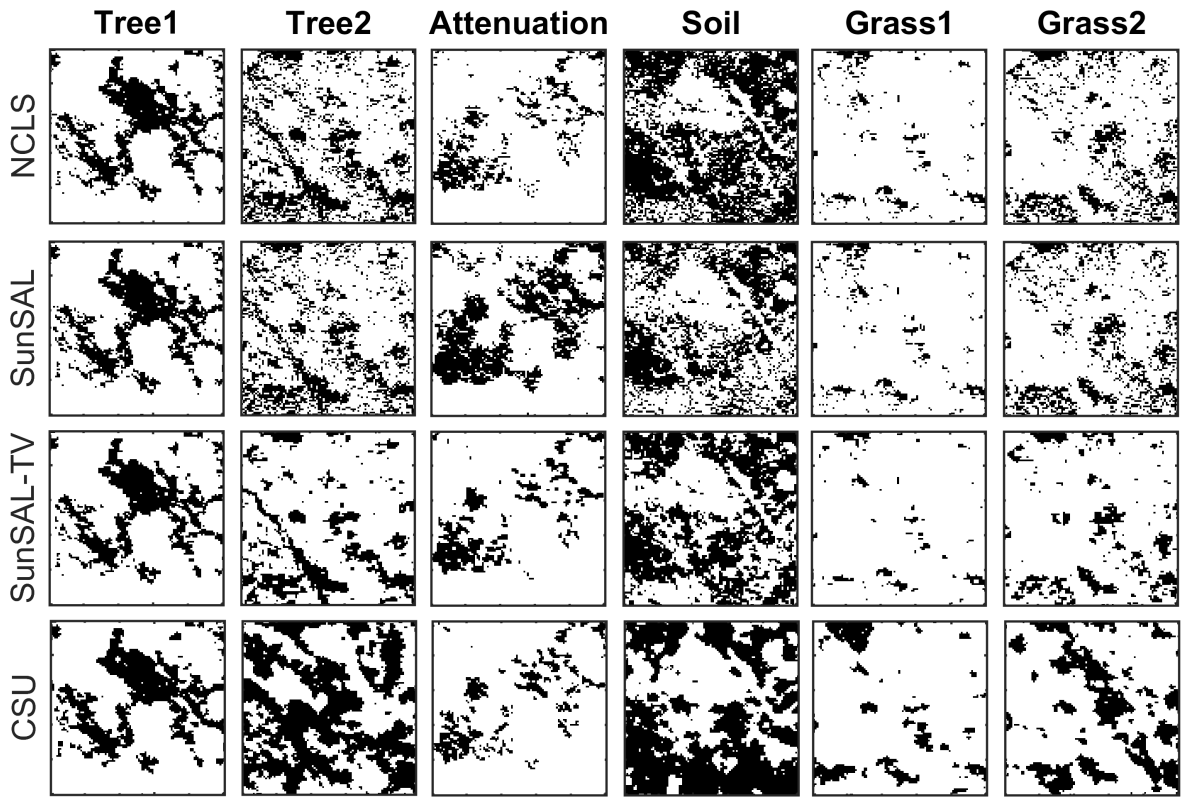

Fig. 7. Estimated presence maps (white pixels correspond to regions where the endmembers are present) for the real Villelongue image.

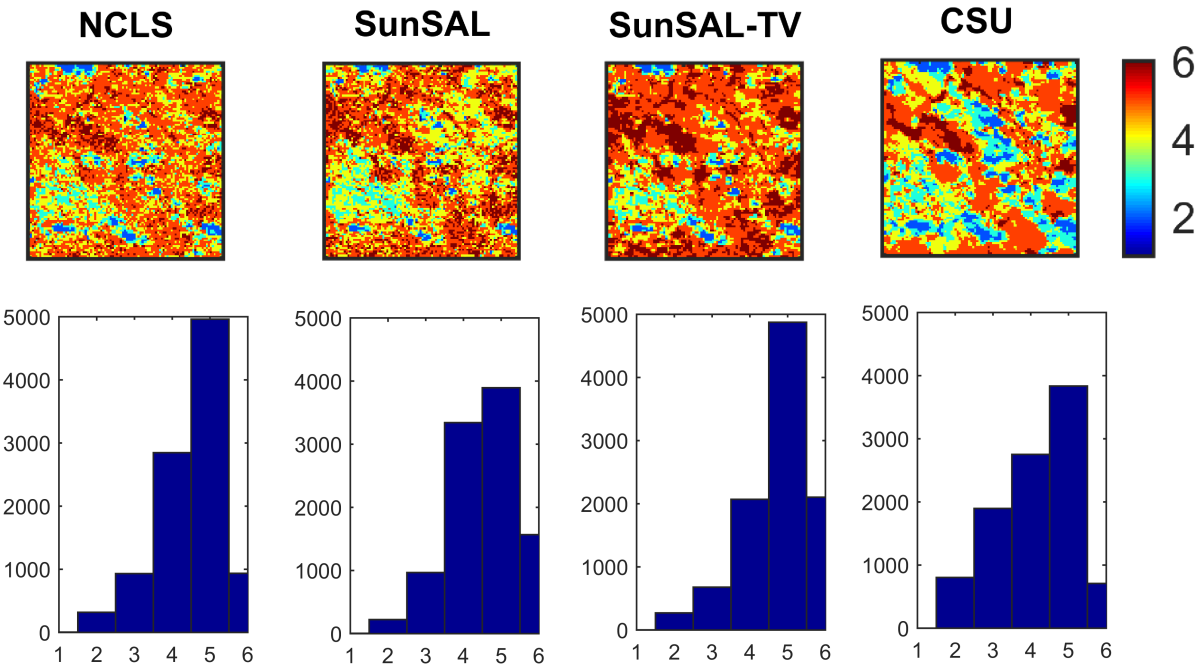

Fig. 8. Top: Estimated number of active endmembers per pixel in the Villelongue scene. Bottom: Histograms of the estimated numbers of active endmembers per pixel (computed over the 10^5 image pixels)

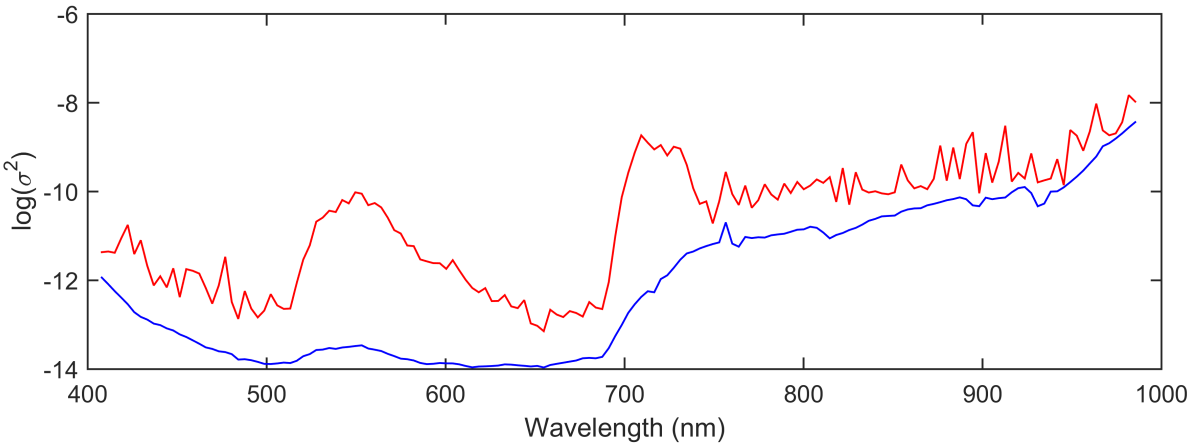

Fig. 9. Noise variances estimated by the CSU (blue) and the Hysime algorithms (red) for the real Villelongue image.

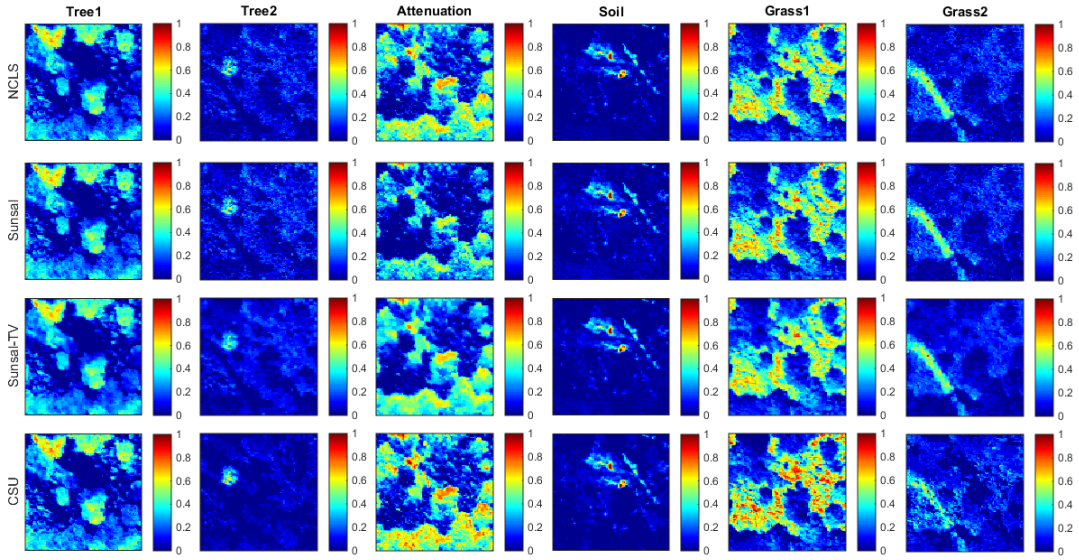

Fig. 10. Estimated abundance maps for the real Villelongue image.

REFERENCES

[1] J. M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, and J. Chanussot, "Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches," *IEEE J. Sel. Topics Appl. Earth Observations Remote Sensing*, vol. 5, no. 2, pp. 354–379, April 2012.

[2] J. Bioucas-Dias, A. Plaza, G. Camps-Valls, P. Scheunders, N. Nasrabadi, and J. Chanussot, "Hyperspectral remote sensing data analysis and future challenges," *Geoscience and Remote Sensing Magazine, IEEE*, vol. 1, no. 2, pp. 6–36, June 2013.

[3] N. Dobigeon, J.-Y. Tourneret, C. Richard, J. C. M. Bermudez, S. McLaughlin, and A. O. Hero, "Nonlinear unmixing of hyperspectral images: Models and algorithms," *IEEE Signal Processing Magazine*, vol. 31, no. 1, pp. 82–94, Jan. 2014.

[4] M. Craig, "Minimum volume transforms for remotely sensed data," *IEEE Trans. Geosci. and Remote Sensing*, vol. 32, no. 3, pp. 542–552, May 1994.

[5] D. C. Heinz and C.-I Chang, "Fully constrained least-squares linear spectral mixture analysis method for material quantification in hyperspectral imagery," *IEEE Trans. Geosci. and Remote Sensing*, vol. 29, no. 3, pp. 529–545, March 2001.

[6] O. Eches, N. Dobigeon, C. Mailhes, and J.-Y. Tourneret, "Bayesian estimation of linear mixtures using the normal compositional model," *IEEE Trans. Image Processing*, vol. 19, no. 6, pp. 1403–1413, June 2010.

[7] L. Miao, H. Qi, and H. Szu, "A maximum entropy approach to unsupervised mixed-pixel decomposition," *IEEE Trans. Image Processing*, vol. 16, no. 4, pp. 1008–1021, april 2007.

[8] Z. Yang, G. Zhou, S. Xie, S. Ding, J.-M. Yang, and J. Zhang, "Blind spectral unmixing based on sparse nonnegative matrix factorization," *IEEE Trans. Image Processing*, vol. 20, no. 4, pp. 1112–1125, April 2011.

[9] M.-D. Iordache, J. Bioucas-Dias, and J. Nascimento, "Sparse unmixing of hyperspectral images," *IEEE Trans. Geosci. and Remote Sensing*, vol. 49, no. 6, pp. 2014–2039, June 2011.

[10] J. Bioucas-Dias and M. A. T. Figueiredo, "Alternating direction algorithms for constrained sparse regression: application to hyperspectral unmixing," in *Proc. IEEE GRSS Workshop Hyperspectral Image SIgnal Process.: Evolution in Remote Sens. (WHISPERS)*, Reykjavík, Iceland, 2010, pp. 1–4.

[11] M.-D. Iordache, J. Bioucas-Dias, and A. Plaza, "Collaborative sparse regression for hyperspectral unmixing," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 52, no. 1, pp. 341–354, Jan 2014.

[12] M.-D. Iordache, J. Bioucas-Dias, A. Plaza, and B. Somers, "Music-csr: Hyperspectral unmixing via multiple signal classification and collaborative sparse regression," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 52, no. 7, pp. 4364–4382, July 2014.

[13] N. Dobigeon, J.-Y. Tourneret, and C.-I Chang, "Semi-supervised linear spectral unmixing using a hierarchical Bayesian model for hyperspectral imagery," *IEEE Trans. Signal Process.*, vol. 56, no. 7, pp. 2684–2695, July 2008.

[14] O. Eches, N. Dobigeon, and J.-Y. Tourneret, "Estimating the number of endmembers in hyperspectral images using the normal compositional model an a hierarchical bayesian algorithm," *IEEE J. Sel. Topics Signal Processing*, vol. 3, no. 3, pp. 582–591, June 2010.

[15] L. Miao and H. Qi, "Endmember extraction from highly mixed data using minimum volume constrained nonnegative matrix factorization," *IEEE Trans. Geosci. and Remote Sensing*, vol. 45, no. 3, pp. 765–777, March 2007.

[16] J. Bioucas-Dias, "A variable splitting augmented Lagrangian approach to linear spectral unmixing," in *Proc. IEEE GRSS Workshop Hyperspectral Image SIgnal Process.: Evolution in Remote Sens. (WHISPERS)*, Grenoble, France, 2009, pp. 1–4.

[17] J. Nascimento and J. Bioucas-Dias, "Hyperspectral unmixing based on mixtures of dirichlet components," *IEEE Trans. Geosci. and Remote Sensing*, vol. 50, no. 3, pp. 863–878, March 2012.

[18] O. Eches, N. Dobigeon, and J.-Y. Tourneret, "Enhancing hyperspectral image unmixing with spatial correlations," *IEEE Trans. Geosci. and Remote Sensing*, vol. 49, no. 11, pp. 4239–4247, Nov. 2011.

[19] M.-D. Iordache, J. Bioucas-Dias, and A. Plaza, "Total variation spatial regularization for sparse hyperspectral unmixing," *IEEE Trans. Geosci. and Remote Sensing*, vol. 50, no. 11, pp. 4484–4502, Nov 2012.

[20] O. Eches, J. Benediktsson, N. Dobigeon, and J.-Y. Tourneret, "Adaptive markov random fields for joint unmixing and segmentation of hyperspectral images," *IEEE Trans. Image Processing*, vol. 22, no. 1, pp. 5–16, Jan 2013.

[21] J. Chen, C. Richard, and P. Honeine, "Nonlinear estimation of material abundances in hyperspectral images with ell_1-norm spatial regularization," *IEEE Trans. Geosci. and Remote Sensing*, vol. 52, no. 5, pp. 2654–2665, May 2014.

[22] Y. Altmann, N. Dobigeon, S. McLaughlin, and J.-Y. Tourneret, "Residual component analysis of hyperspectral images - application to joint nonlinear unmixing and nonlinearity detection," *IEEE Trans. Image Processing*, vol. 23, no. 5, pp. 2148–2158, May 2014.

[23] G. E. Newstadt, A. O. H. III, and J. Simmons, "Robust spectral unmixing for anomaly detection," in *Proc. IEEE-SP Workshop Stat. and Signal Processing*, Golden Coast, July 2014.

[24] J. M. Bernardo and A. F. M. Smith, *Bayesian Theory*. New York: John Wiley & Sons, 1994.

[25] M. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tourneret, "Estimating the granularity coefficient of a Potts-Markov random field within an MCMC algorithm," *IEEE Trans. Image Processing*, vol. 22, no. 6, pp. 2385–2397, June 2013.

[26] M. Pereyra, N. Whiteley, C. Andrieu, and J.-Y. Tourneret, "Maximum marginal likelihood estimation of the granularity coefficient of a Potts-Markov random field with an MCMC algorithm," in *Proc. IEEE-SP Workshop Stat. and Signal Processing*, Golden Coast, July 2014.

[27] C. P. Robert, *The Bayesian Choice: from Decision-Theoretic Motivations to Computational Implementation*, 2nd ed., ser. Springer Texts in Statistics. New York: Springer-Verlag, 2007.

[28] A. Doucet, S. J. Godsill, and C. P. Robert, "Marginal maximum a posteriori estimation using markov chain monte carlo," *Statistics and Computing*, vol. 12, no. 1, pp. 77–84, 2002.

[29] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed. New York: Springer-Verlag, 2004.

[30] A. Pakman and L. Paninski, "Exact Hamiltonian Monte Carlo for Truncated Multivariate Gaussians," *ArXiv e-prints*, Aug. 2012.

[31] R. Clark, G. Swayze, R. Wise, E. Livo, T. Hoefen, R. Kokaly, and S. Sutley, "Usgs digital spectral library splib06a: U.s. geological survey," Digital Data Series 231, 2007. [Online]. Available: http://speclab.cr.usgs.gov/spectral.lib06

[32] D. Donoho, E. M., and V. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 50, no. 1, pp. 6–18, Jan. 2006.

[33] D. Sheeren, M. Fauvel, S. Ladet, A. Jacquin, G. Bertoni, and A. Gibon, "Mapping ash tree colonization in an agricultural mountain landscape: Investigating the potential of hyperspectral imagery," in *Proc. IEEE Int. Conf. Geosci. and Remote Sensing (IGARSS)*, July 2011, pp. 3672–3675.

[34] Y. Altmann, N. Dobigeon, S. McLaughlin, and J. Tourneret, "Nonlinear spectral unmixing of hyperspectral images using Gaussian processes," *IEEE Trans. Signal Process.*, vol. 61, no. 10, pp. 2442–2453, May 2013.

[35] J. M. Bioucas-Dias and J. M. P. Nascimento, "Hyperspectral subspace identification," *IEEE Trans. Geosci. and Remote Sensing*, vol. 46, no. 8, pp. 2435–2445, Aug. 2008.